On (Commercial) Benefits of Automatic Text Summarization Systems in the News Domain: A Case of Media Monitoring and Media Response Analysis

Pashutan Modaresi and **Philipp Gross** and **Siavash Sefidrodi** and **Mirja Eckhof**
pressrelations GmbH
D-40211 Düsseldorf, Germany
[first name].[last name]@pressrelations.de

Stefan Conrad
Institute of Computer Science
Heinrich Heine University of Düsseldorf
conrad@cs.uni-duesseldorf.de

arXiv:1701.00728v1 [cs.CL] 3 Jan 2017

Abstract

In this work, we present the results of a systematic study to investigate the (commercial) benefits of automatic text summarization systems in a real world scenario. More specifically, we define a use case in the context of media monitoring and media response analysis and claim that even using a simple query-based extractive approach can dramatically save the processing time of the employees without significantly reducing the quality of their work.

1 Introduction

Automatic text summarization has been an evolving field of research. Having started with the pioneering work of Luhn (Luhn, 1958), specifically in recent years, automatic text summarization has made remarkable signs of progress with the popularity of deep learning approaches (Rush et al., 2015; Chopra et al., 2016).

Providing a formal definition of automatic text summarization is rather a challenging task. This work pursues the following definition: Given a set Q of queries, automatic text summarization is a reductive transformation of a collection of documents D with $|D| > 0$ into a single or multiple target document(s), where the target document(s) are more readable than the documents in D and contain the relevant information of D according to Q (Modaresi and Conrad, 2015). This definition, comprises both extractive and abstractive approaches, where by extractive we mean methods that select the most salient sentences in a document and by abstractive we mean methods that incorporate language generation to reformulate a document in a reductive way.

Automatic text summarization has been applied to many domains, among which is the news domain the focus of this work. Despite many attempts to improve the performance of summarization systems (Ferreira et al., 2016; Wei and Gao, 2015), to the best knowledge of the authors, no systematic study was performed to investigate the (commercial) benefits of the summarization systems in a real world scenario.

We claim that using (even very) simple automatic summarization systems can dramatically improve the workflow of employees without affecting their quality of work. To investigate our claim we define a use case in the context of media monitoring and media response analysis (Section 2) and establish several criteria to measure the effectiveness of the summarization systems in our use case (Section 3). In Section 4 we discuss the design of our experiment and report the results in Section 5. Finally, in Section 6 we conclude our work.

2 Use Case Definition

We investigate the (commercial) benefits of integrating an automatic summarization system in the semi-automatic workflows of media analysts doing *media monitoring* and *media response analysis* at pressrelations GmbH[1]. In the following, we

[1] http://www.pressrelations.de/

shortly define the terms as mentioned above.

The goal of media monitoring is to gather all relevant information on specific topics, companies or organizations. To this end, search queries are defined, with which the massive amount of available information can be filtered automatically. Typically, in a post-processing step, the quality of the gathered information is increased using manual filtering by trained media analysts.

In media response analysis, the publications in the media (print media, radio, television, and online media) are evaluated according to various predefined criteria. As a result of this, it is possible to deduce whether and how journalists have recorded and processed the PR (Public Relations) messages. Possible questions to be answered in the context of media response analysis are: How are the publications distributed over time? How many listeners, viewers or readers were potentially reached? What are the tonality and opinion tendency of the publications? (Grupe, 2011)

Typically, analysis results are given to the clients in the form of extensive reports. In the case of textual media, the immense amount of time required to read texts and to write abstracts and reports is a high cost factor in the preparation of media resonance analysis reports.

We claim that the described process can be partially optimized by incorporating automatic summarization systems, leading to remarkable financial advantages for the companies.

3 Evaluation Criteria

From the commercial and academic point of view, the *quality* of the summaries plays an import role. Various automatic methods such as ROUGE (Lin, 2004), BLEU (Papineni et al., 2002), and Pyramid (Nenkova et al., 2007) have been used successfully to evaluate the quality of the summaries. Moreover, manual evaluation has also been incorporated for quality assessment of the summaries (Modaresi and Conrad, 2014). Another important criterion that is mostly neglected in academic publications is the *gain in time*, defined as the amount of saved time by a user through the usage of the summaries.

In our use case, the quality of a summary comprises of two aspects: *completeness* and *readability*. The term *completeness*, describes the requirement of a summary to contain all relevant information of an article. The relevance of information is determined based on a *query*. For instance, the query might be a named entity, and we expect that the summary contains all relevant information regarding the named entity.

The term *readability* refers to the coherence and the grammatical correctness of the summary. While the grammatical correctness is defined at the sentence level, the coherence of the summary is determined on the whole text. That means that the sentences of the summary should not only be grammatically correct in isolation, but also they must be coherent to make the summary readable.

Both *completeness* and *readability* are criteria that are difficult to evaluate and define formally, and it has been shown that they are both very subjective criteria, where their assessment varies from person to person (Torres-Moreno, 2014). In the case of completeness, it is unclear how to formalize the relevance of information, and in in the case of readability the same holds for the concept of coherence.

Therefore, we define the quality of a summary from a practical and commercial point of view. For this, we define the quality of a summary in terms of a binary decision problem where the question to be asked is: *can the produced summary in its current form be delivered to a customer or not?*

Furthermore, in our use case, the *gain in time* is defined as the processing time that can be saved by media analysts, assisted by a summarization system. It should be clear that the reduction of the processing time could lead to the reduction of costs in a company.

In the following section, the design of our experiment with respect to the criteria mentioned above (quality and gain in time) will be explained.

4 Experiment Setup

To conduct our experiments we incorporated eight media analysts (specialists in writing summaries for customers) and divided them into two equisized groups. One group received only the news articles (Group A), and the other one received only the query-based extracted summaries without having access to the original articles. Given a query consisting of a single named entity, both groups were asked to write summaries with the following properties:

- The summary should be compact and consist of maximum two sentences.

- The summary should contain the main topic of the article and also the most relevant information regarding the query.

As previously stated, the summaries created by media analysts were evaluated based on two criteria: *quality* and *gain in time*. The *gain in time* was measured automatically using a web interface by tracking the processing time of the media analysis for creating the text summaries. We interpret the *gain in time* as the answer to the question: *In average, what percentage faster/slower is group A in compare to group B?*. Let $\tilde{t}_A$ and $\tilde{t}_B$ be the average processing times of the media analysts in group A and B respectively. We define *gain in time* as in Equation 1.

$$g(\tilde{t}_A, \tilde{t}_B) = \begin{cases} 100 \cdot \left|1 - \frac{\tilde{t}_A}{\tilde{t}_B}\right| & \text{if } \tilde{t}_A \leqslant \tilde{t}_B \\ 100 \cdot \left|1 - \frac{\tilde{t}_B}{\tilde{t}_A}\right| & \text{if } \tilde{t}_A > \tilde{t}_B \end{cases} \quad (1)$$

Notice that it holds $g(\tilde{t}_A, \tilde{t}_B) = g(\tilde{t}_B, \tilde{t}_A)$ and g reflects only the magnitude of the saved time and not its direction. The direction can be determined based on the values of $\tilde{t}_A$ and $\tilde{t}_B$.

On the other hand, the quality of the summaries was evaluated by a curator (an experienced media analyst in direct contact with customers). The curator was responsible for evaluating the summaries created by media analysts in both groups and scored them with a zero or a one. With zero meaning that the quality of the summary is not sufficient and the product cannot be delivered to the client and with one meaning the quality of the summary is sufficient enough to be delivered to the customer. Let the vector q of size m be a one-hot vector consisting of 0s and 1s, where the i-th element in q represents the evaluation of the curator for the i-th summary among the m available summaries. Given that, we compute the average summary quality of a set of summaries by computing the average of its corresponding evaluation vector q.

In total, ten news articles were provided to the media analysts. The articles for group A had an average word count of 1438 with the standard deviation being 497. Group B received only the summaries of the articles, created automatically with a heuristic-based approach. The automatically generated summaries had an average length of 81 words with the standard deviation being 23.

Algorithm 1 Query-based Summarization

```
1:  procedure SUMMARIZE(T, Q)
2:      S ← ∅
3:      T' ← Segment(T)
4:      E ← EntityDistribution(T')
5:      m ← Median(E)
6:      E' ← ∅
7:      for e in E do
8:          if freq(e) > m then
9:              E' ← E' ∪ e
10:     S ← Lead(T')
11:     S ← QueryMatch(T', Q)
12:     S ← CentralEntityMatch(T', E')
13:     return S
```

The pseudocode of the invoked query-based extractive summarizer is depicted in Algorithm 1.

In line 2 the summary S is initialized with an empty set. Given the input text T, the text is segmented into sentences and stored in the list T' (line 2). In line 3, the named entities of the text are recognized and stored in a dictionary where each key represent a named entity, and its corresponding value is the frequency of the named entity in the text. Lines 5-9 depict the procedure to select *central named entities*. Let m be the median of the named entities frequencies. A named entity e is called a *central named entity* if its frequency in the text is higher than the twice of the median. In line 10 we add the lead of the news article to the summary, as the lead usually can be interpreted as a compact summary of the whole article. Afterwards in line 11, the sentences that contain the query Q are added to the summary. Finally, we extend the list of summary sentences with sentences containing the *central named entities* and return the summary.

5 Results

In total, we collected 80 summaries created by the media analysts in both groups. For each summary, its processing time and its quality evaluated by a curator was recorded. Based on the collected data, we answered the following questions:

1. Intergroup processing time: Is there a significant difference between the processing times of individual media analysts in a group?

2. Intergroup quality: Is there a significant difference between the quality of the created summaries by the media analysts in a group?

3. Intragroup processing time: Is there a significant different between the average processing times of media analysts in groups A and B? If so, which group has a faster processing time?

4. Intragroup quality: Is there a significant difference between the average qualities of created summaries by media analysts in groups A and B? If so, which group created more qualitative summaries?

The remaining of this section reports the answers to the above questions.

5.1 Intergroup Processing Time

The processing times of the media analysts in group A (A1-A4) and group B (B1-B4) are visualized using boxplots in Figures 1a and 1b respectively. In both groups, the differences among the average processing times are observable. Our goal is to investigate whether the differences between the processing times of media analysts is statistically significant.

To compare the means of processing times among the media analysts in a group we use the *one-way analysis of variance* (one-way ANOVA). The null hypothesis in the ANOVA test is that the mean processing times of the media analysts in a group are the same. To perform the ANOVA test we first examine if the requirements of the ANOVA test are satisfied (Miller, 1997).

The first requirement of the ANOVA test is that the processing times of the individual media analysts are normally distributed. For this, we use the Shapiro-Wilk test (Shapiro and Wilk, 1965) with the null hypothesis being that the processing times are normally distributed. Table 1 reports the results of the test.

Media Analyst	W	p-value
A1	0.90244	0.233
A2	0.91638	0.3277
A3	0.76592	0.0055
A4	0.73605	0.0024
B1	0.85143	0.0604
B2	0.95625	0.7425
B3	0.94609	0.6226
B4	0.93536	0.5026

Table 1: Shapiro-Wilk Test for Processing Times

In Table 1, W is the test statistic and we reject the null hypothesis if the p-value is less than the chosen significance level $\alpha = 0.05$. Thus the null hypothesis will be rejected for A3, A4, and B1, meaning that the processing times of them are not normally distributed. For other media analysts, the normality assumption holds. Although in several cases the normality requirement of the ANOVA test is violated, it is still possible to use the ANOVA test, as it was shown that the ANOVA test is relatively robust to the violation of the normality requirement (Kirk, 2012).

The second requirement to perform the ANOVA test is that the processing times of the media analysts have equal variances. For this, we use the Bartlett's test (Dalgaard, 2008) with the null hypothesis that the processing times of the media analysts have the same variance. The results of the Bartlett's test for groups A and B are reported in Table 2

Group	χ^2	p-value
A	4.3726	0.2239
B	6.9013	0.0751

Table 2: Bartlett's Test for Processing Times

In Table 2, χ^2 is the test statistic and we reject the null hypothesis if the p-value is less than the chosen significance level $\alpha = 0.05$. For both groups, the p-value is greater than the significance level, and thus there is no evidence that the variances of processing times of individual media analysts are different.

Having investigated the assumptions of the ANOVA test, we now report the results of the ANOVA test (See Table 3).

Group	F value	p-value
A	$1.413 \cdot 10^{33}$	$< 2 \cdot 10^{-16}$
B	$4.2 \cdot 10^{34}$	$< 2 \cdot 10^{-16}$

Table 3: ANOVA Test for Processing Times

In Table 3, the F value is the F test statistic and we reject the null hypothesis if the p-value is less than the chosen alpha level $\alpha = 0.05$. Thus, the mean processing times of media analysts in group A are not the same and there is a significant difference between them. The same hold for group B.

The so far shown results crystallize an important property of the summarization process. Given the same set of news articles and the same brief-

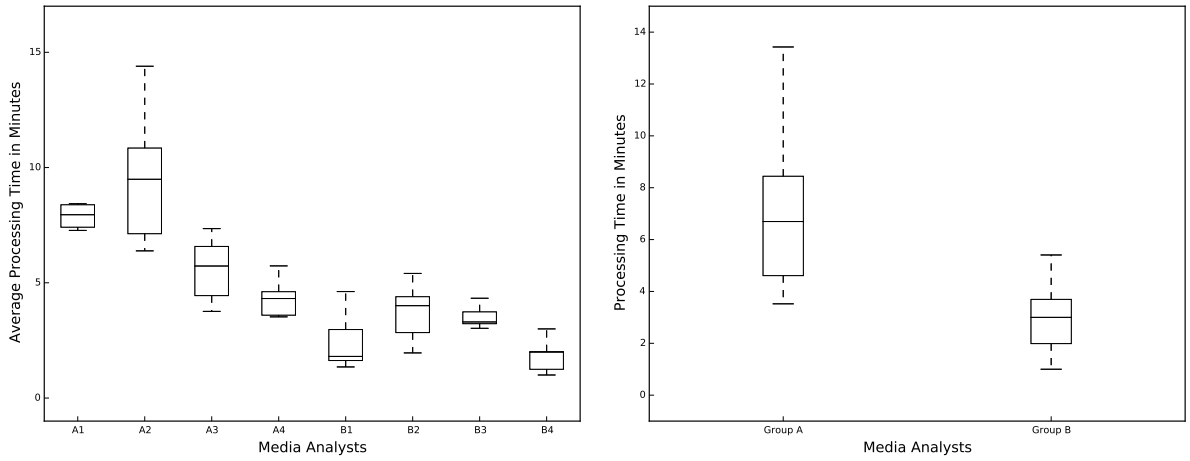

(a) Processing Times of Individual Media Analysts

(b) Processing Times of Groups A and B

Figure 1: Comparison of the Processing Times

ing to all media analysts, the average time required by the media analysts within a group to summarize the articles is significantly different from each other.

5.2 Intergroup Quality

The results of the manual evaluation of the summaries by the curator are represented in Table 4. In this section, our goal is to systematically investigate whether the qualities of the summaries produced by media analysts in a group are significantly different from each other.

Different from the previous section where we compared the processing times of the media analysts in a group using the ANOVA test, the comparison of the qualities among the media analysts cannot be performed using the ANOVA test (due to the huge violation of the normality assumption). Therefore, we interpret the evaluation results of each media analyst as a Binomial distribution $\mathrm{B}(n, p)$ with $n = 10$ (number of articles shown to each media analyst) and p being the numbers of times the curator was satisfied with the quality of the summaries created by the media analyst.

	Group A				Group B			
	A1	A2	A3	A4	B1	B2	B3	B4
Quality	0.9	1.0	0.8	0.6	0.3	0.9	1.0	0.7
Overall	0.82				0.72			

Table 4: Results of Manual Evaluation of Quality

To test whether the qualities of the produced summaries are significantly different from each other, we use the Fisher's Exact Test (Dalgaard, 2008) with the null hypothesis that the qualities are not different from each other. For group A we have a p-value of 0.1087, and for group B the p-value is 0.0022. Thus the null hypothesis can only be rejected for group B. The so far shown results lead us to the following conclusions: Given the news articles, no significant difference among the qualities of the produced summaries by the media analysts can be observed. Furthermore, given only the automatically created summaries, the media analysts produce summaries with significantly different qualities.

5.3 Intragroup Processing Time

So far, we only investigated the intergroup properties. In this section, we answer the question whether there exists a significant difference between the average processing times of group A and B?

In Figure 1b, the processing times of the groups A and B are compared using boxplots. Using the Equation 1, we compute the *gain in time* for group B, that is roughly 58%, meaning that as expected, the media analysts in group B required much less time to create the summaries in compare to the media analysts in group A. Similar to the Section 5.1, we use the ANOVA test to check the significance of this outcome. The results of the test are reported in Table 5.

In Table 5, F value is the F test statistic and we reject the null hypothesis if the p-value is less than the chosen alpha level $\alpha = 0.05$. Thus, the processing times of media analysts in group B are sig-

Group	F value	p-value
A vs. B	$2.14 \cdot 10^{33}$	$< 2 \cdot 2^{-16}$

Table 5: ANOVA Test for Intragroup Processing Times

nificantly lower than the processing times of media analysts in group A.

The results show that using a simple query-based extractive summarization system, the media analysts had a significant gain in time by the process of creating the text summaries.

5.4 Intragroup Quality

In the final step, we compare the quality of the produced summaries between both groups and answer the question whether there is a significant difference between the qualities? To answer this question we perform the Fisher's Exact Test and obtain the p-value of 0.4225. Thus the null hypothesis of the test cannot be rejected and we conclude that the qualities of the summaries among both groups are not significantly different.

Using the results above, we conclude that providing the media analysts with automatically created summaries does not have a negative impact on the quality of the summaries they generated and no significant difference in quality could be observed in compare to the media analysts that had access to the full new articles.

6 Conclusions

To investigate the (commercial) benefits of the summarization systems, we designed an experiment where two groups of media analysts were given the task to summarize news articles. Group A received the whole news articles and group B received only the automatically created text summaries. In summary, we showed that:

- The average time required by the media analysts within a group to summarize the articles is significantly different from each other.
- Given the news articles, no significant difference among the qualities of the produced summaries by the media analysts can be observed. Furthermore, given only the automatically created summaries, the media analysts produce summaries with significantly different qualities.
- The media analysts had a significant gain in time by the process of creating the text summaries (58%).
- Providing the media analysts with automatically created summaries does not have a negative impact on the quality of the summaries they generated

The results mentioned above indicate that incorporating even simple summarization systems can dramatically improve the workflow of the employees.

For future work we plan to repeat our experiment with more sophisticated summarization algorithms and compare the *gain in time* to our baseline setting. Furthermore, we plan to increase the number of media analysts to obtain more reliable results.

Acknowledgments

This work was funded by the German Federal Ministry of Economics and Technology under the ZIM program (Grant No. KF2846504).

References

[Chopra et al.2016] Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. Abstractive Sentence Summarization with Attentive Recurrent Neural Networks. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 93–98.

[Dalgaard2008] P. Dalgaard. 2008. *Introductory Statistics with R*. Statistics and Computing. Springer New York.

[Ferreira et al.2016] Rodolfo Ferreira, Rafael Ferreira, Rafael Dueire Lins, Hilário Oliveira, Marcelo Riss, and Steven J. Simske. 2016. Applying Link Target Identification and Content Extraction to Improve Web News Summarization. In *Proceedings of the 2016 ACM Symposium on Document Engineering*, DocEng '16, pages 197–200. ACM.

[Grupe2011] S. Grupe. 2011. *Public Relations: Ein Wegweiser für die PR-Praxis*. Springer Berlin Heidelberg.

[Kirk2012] R.E. Kirk. 2012. *Experimental Design: Procedures for the Behavioral Sciences: Procedures for the Behavioral Sciences*. SAGE Publications.

[Lin2004] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81. Association for Computational Linguistics.

[Luhn1958] H. P. Luhn. 1958. The Automatic Creation of Literature Abstracts. *IBM J. Res. Dev.*, 2(2):159–165, April.

[Miller1997] R.G. Miller. 1997. *Beyond ANOVA: Basics of Applied Statistics*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.

[Modaresi and Conrad2014] Pashutan Modaresi and Stefan Conrad. 2014. From Phrases to Keyphrases: An Unsupervised Fuzzy Set Approach to Summarize News Articles. In *Proceedings of the 12th International Conference on Advances in Mobile Computing and Multimedia*, pages 336–341.

[Modaresi and Conrad2015] Pashutan Modaresi and Stefan Conrad. 2015. On Definition of Automatic Text Summarization. In *Proceedings of Second International Conference on Digital Information Processing, Data Mining, and Wireless Communications*, DIPDMWC2015, pages 33–40. SDIWC.

[Nenkova et al.2007] Ani Nenkova, Rebecca Passonneau, and Kathleen McKeown. 2007. The Pyramid Method: Incorporating Human Content Selection Variation in Summarization Evaluation. *ACM Trans. Speech Lang. Process.*, 4(2).

[Papineni et al.2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pages 311–318.

[Rush et al.2015] Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A Neural Attention Model for Abstractive Sentence Summarization. *CoRR*, abs/1509.00685.

[Shapiro and Wilk1965] S. S. Shapiro and M. B. Wilk. 1965. An Analysis of Variance Test for Normality (Complete Samples). *Biometrika*, 52(3/4):591–611.

[Torres-Moreno2014] Juan-Manuel Torres-Moreno. 2014. *Automatic Text Summarization*. Wiley.

[Wei and Gao2015] Zhongyu Wei and Wei Gao. 2015. Gibberish, Assistant, or Master?: Using Tweets Linking to News for Extractive Single-Document Summarization. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 1003–1006. ACM.